• The solutions provided hereon are for reference and study tools and are not to be considered final exam-assessment criteria

1. Explain the operation of the traceroute Linux utility: the underlying protocol stack it uses for exchanging messages with the routers, the relevant fields of the exchanged messages and, also, what it serves for. In this question it's important that you use as rich and accurate a terminology as possible.

The traceroute utility in Linux, sent from an IP host, estimates the path to an Internet destination by sending UDP/IP messages with increasing TTL. The first message will contain TTL=1 which will cause the first router visited by the packet to send an ICMP TTL Time Exceeded message (TEM) back to the sending host. Ensuing UDP/IP messages will follow, each with a larger value of TTL which will elicit the ICMP TEM from routers farther away from the sending host. Each TEM received allows the sending host to record a new router on the source-to-destination path and the estimated Rtt (Round Trip Time).

Figure 1 represents the protocol stack resulting on the sending host (On the Linux Op Sys). The router that obtains zero from the decrement of the TTL field of the received UDP packet reacts by sending back an ICMP TEM which protocol stack is depicted on Figure 2. The Ethernets appearing on the bottom-most layer, are obviously an illustrating example out of many other possible Subnetwork-layer physical transmission protocols available (Wi-Fi, PPP, PPoE, etc.).

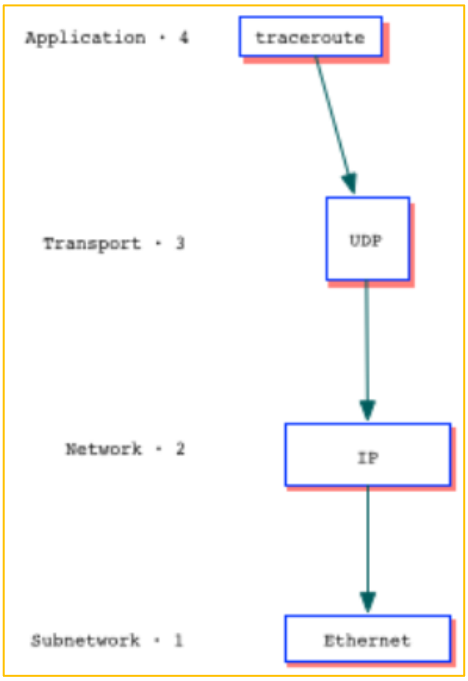

Figure 1. traceroute utility, UDP TTL-probe protocol stack

The essential fields of the involved PDUs are:

- UDP Probe: Any unassigned SRC and DST port numbers
- IP packet containing UDP probe: DST IP, SRC IP, TTL (Time To Live)
- ICMP Time Exceeded Message (TEM): Type and Code fields, sequence numbers
- IP packet containing ICMP TEM: DST IP = Source host, SRC IP = Responding router's IP, Protocol = 1

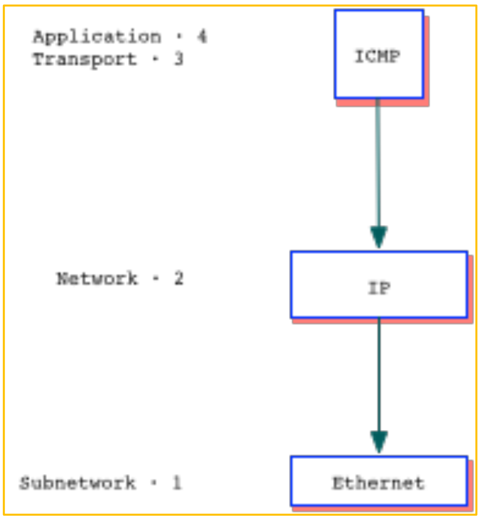

Figure 2. Protocol stack of a router sending back an ICMP TEM message

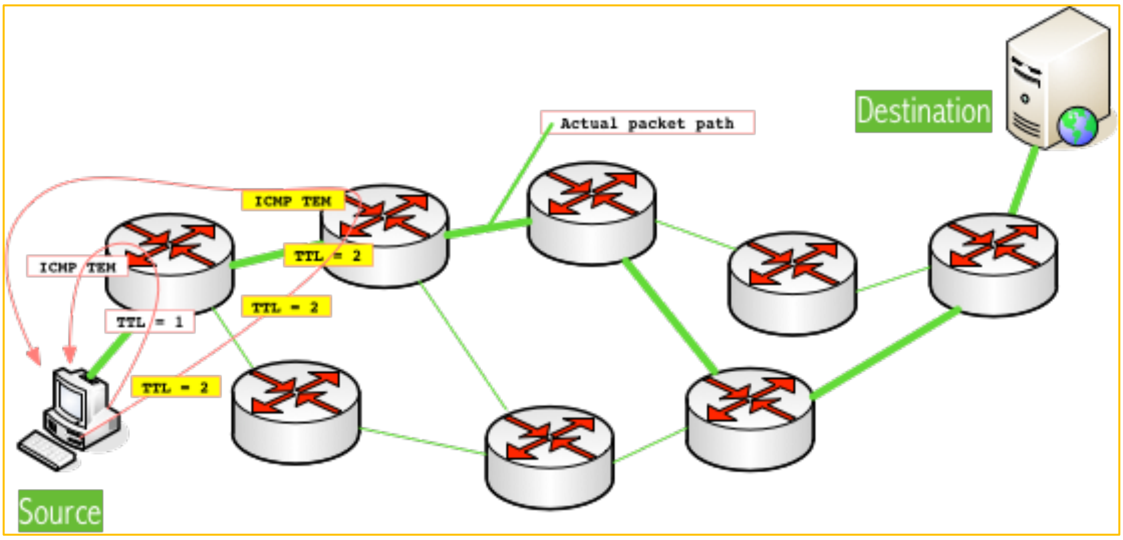

Figure 3. Source host sending UDP/IP packets with increasing TTL values and eliciting receiving routers to send back ICMP TEM messages

2. Calculate how much time it takes to transfer a file which size is 1473 Bytes from a host A to a host B assuming a direct connection between them. The Rtt is 20 µs and the transmission speed is 10 Mbps. Assume that the sending application at host A uses the following protocol stack: UDP/IP/Ethernet. The involved header sizes follow:

   - Size of a UDP datagram header is ~~20~~ <u>8</u> Bytes
   - Size of an IP packet header is 20 Bytes
   - Ethernet Datalink header is comprised of:
     - Dest. Address (DST MAC): 6 Bytes
     - Src. Address (SRC MAC): 6 Bytes
     - Ethertype: 2 Bytes
   - The Max payload size that can be encapsulated into an Ethernet frame is 1500 Bytes
   - The Physical-layer Ethernet header has two sections:
     - Preamble: 64 bits
     - CRC: 32 bits

The problem statement requests that we calculate the time that it takes transferring the full 1473-Byte file from host A to host B assuming both hosts are directly connected at an unspecified distance (See Figure 1, below). Since distance D affects the total transfer time (By T_p, the propagation time, which is equal to $\frac{Rtt}{2}$), in the total transfer time we express T_p as a function of D).

In general, UDP's socket interface will accept write sizes larger than 1473 bytes, but that entirely depends on the implementation and on certain technical details about the IP configuration such as whether PMTUD is in use or not. We'll assume the following context for responding to the question:

- The MTU is 1500 bytes according to the problem statement which leads us to be sure that the system will fragment any transmitted IP packet larger than that size.
- The system is not doing any kind of Path MTU discovery.

We can expect the datagram socket to accept the 1473-byte sized block of data in a single write operation. UDP will build a datagram containing the 1473 bytes, its total size will be 1473 bytes + 8 bytes (UDP header) = 1481 bytes. Can this size be sent in a single IP packet? An IP packet must fit the MTU, then we ask which is the maximum payload that can be encapsulated into one such IP packet.

Since a *normal (No options)* IP packet has a 20-byte header, the max payload it can carry is = 1500 bytes (MTU) – 20 bytes (IP Header) = 1480 bytes. The UDP datagram that is to be sent weighs 1481 bytes, therefore, IP will fragment it into two IP fragments:

Fragment 1: 20 bytes (Header) + 1480 bytes (Payload)
Fragment 2: 20 bytes (Header) + 1 bytes (Payload)

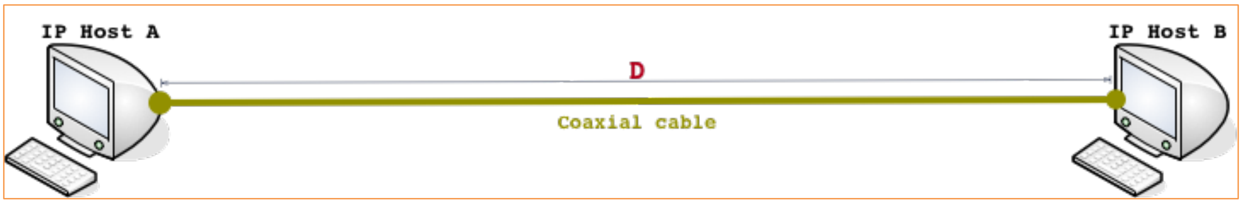

Figure 1. Hosts A and B directly connected at distance D

All in all, the file transfer entails sending two Ethernet frames --probably consecutively, but not necessarily so.

Furthermore, if these two frames were actually sent consecutively, the Ethernet protocol requires that in between the two the sender wait a period of time known as IFG (Inter-frame Gap) which standard time length is 12.5μs. We introduce this technical requirement of Ethernet LANs now, however, since we'll formally introduce it when we take up Ethernet, we will insert no IFG in the solution proposed below (Observe the light blue notice on Figure 2).

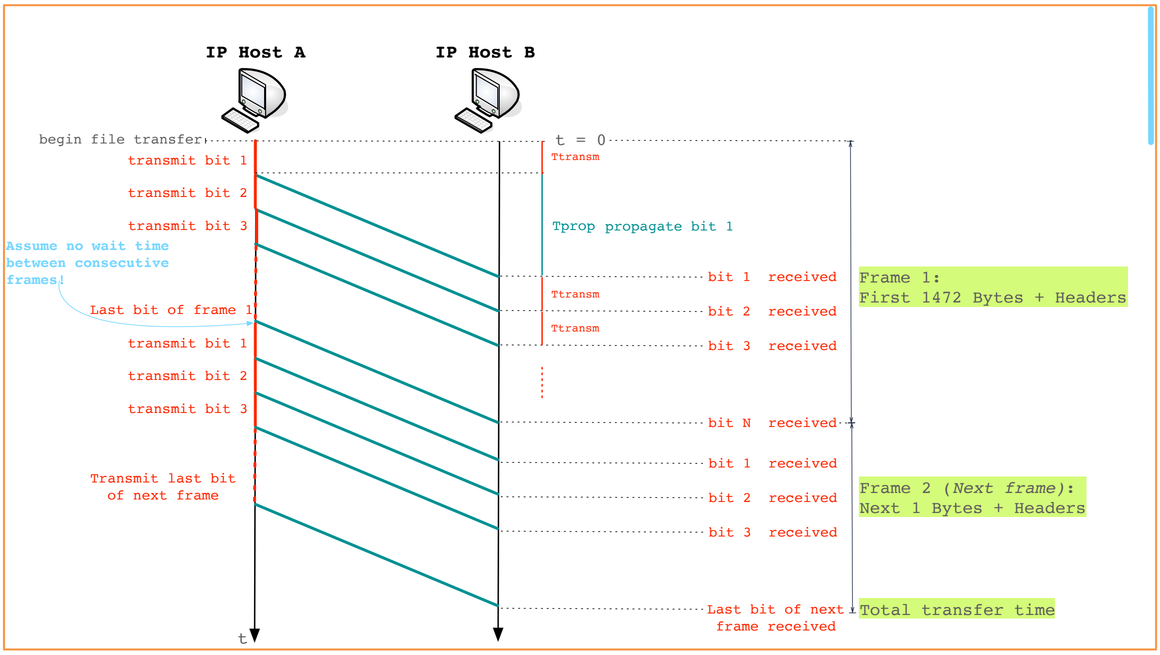

Figure 2. Total transfer time of two back-to-back frames. (In reality, these two frames won't necessarily be transmitted consecutively, but we assume so to enrich this exercise a little).

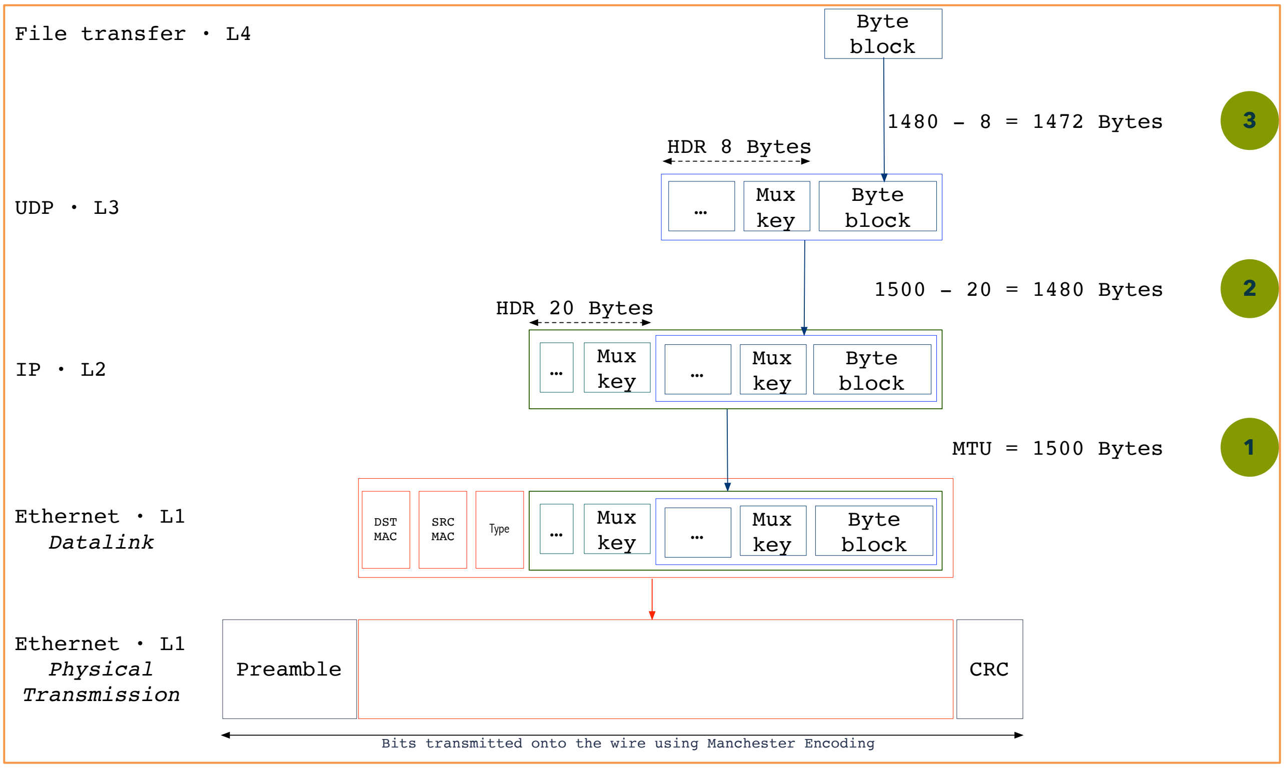

Figure 3. Protocol stack of file application over UDP/IP/Ethernet

By observing the conceptual chronogram in Figure 2, we calculate the total number of bits transmitted (On the wire), assuming the two frames were sent back-to-back and **with no IFG**. The multiplexing hierarchy is developed in Figure 3.

1. **Frame 1/2 (1472 B of payload)**

$$N_{bit-frame\,1} = 1472\,Byte + UDP\ hdr + IP\ hdr + Ethernet\ Datalink\ hdr + Ethernet\ Physical\ hdr$$

$$N_{bit-frame\,1} = 1472\,B \cdot 8\frac{b}{B} + 8B \cdot 8\frac{b}{B} + 20B \cdot 8\frac{b}{B} + 6B \cdot 8\frac{b}{B} + 6B \cdot 8\frac{b}{B} + 2B \cdot 8\frac{b}{B} + 64\,b + 32\,b$$

$$N_{bit-frame\,1} = 11776\,b + 64\,b + 160\,b + 48\,b + 48\,b + 16\,b + 64\,b + 32\,b$$

$$N_{bit-frame\,1} = 12208\,b$$

- The transmission time is the inverse of the transmission speed (100 Mbps): $T_t = \frac{1}{10Mbps} = \frac{1}{10\cdot 10^6 \frac{1}{s}} = 0{,}1 \cdot 10^{-6}s = 100 \cdot 10^{-9}s = 100\,ns$
- If D = 1500m/2 = 750 m, the propagation time is half the Rtt: $T_p = \frac{Rtt}{2} = \frac{51{,}2\,\mu s}{2} = 25{,}6\,\mu s$

2. **Frame 2/2 (1 B of payload)**

$$N_{bit-frame\,2} = 1\,Byte + IP\ hdr + Ethernet\ Datalink\ hdr + Ethernet\ Physical\ hdr$$

$$N_{bit-frame\,2} = 1\,B \cdot 8\frac{b}{B} + 20B \cdot 8\frac{b}{B} + 6B \cdot 8\frac{b}{B} + 6B \cdot 8\frac{b}{B} + 2B \cdot 8\frac{b}{B} + 64\,b + 32\,b$$

$$N_{bit-frame\,2} = 376\,b$$

- Clarification about the propagation time: Since we are **assuming** the two frames are transmitted consecutively and **with IFG = 0 s**, we add no new T_p since it remains concurrent with the transmission time of the ensuing bits, just as happens in the transmission of frame 1.

- The transfer of the two back-to-back frames takes a final time:

$$Total\ transfer\ time = N_{bit-TOTAL} \cdot T_t + T_p = (N_{bit-frame\,1} + N_{bit-frame\,2}) \cdot T_t + T_p$$

$$Total\ transfer\ time = (12208 + 376)b \cdot 100\,\frac{ns}{b} + D \cdot \frac{51.2\,\mu s}{2500\,m}$$

$$\boldsymbol{Total\ transfer\ time = 1258.4\,ms + D \cdot \frac{51.2\,\mu s}{2500\,m}}$$

3. Cite an Internet Application that is elastic from the standpoint of application requirements. Explain the meaning of the word elastic in this context.

Consult slide no. 4 “Cooperating host applications” at chapter 1 section 1 presentation (paloalto.unileon.es/cn)

4. What is the essential identification of hosts in Internet which also serves for locating them?

IP addresses

5. You are part of a team of systems engineers chartered with designing a new NIC (Network Interface Card), specifically you are in charge of integrating the NIC in the Linux operating system and making sure that the current protocols and applications run smoothly over the new NIC. Do you think that the protocols currently running in Linux should be rewritten in order for them to be compatible with the new NIC? What about the LIBPCAP library, should it be adapted to you new NIC?

Current protocols (IP, etc) wouldn’t have to be rewritten whatsoever to make them compatible with the NIC if its Datalink layer exports a consistent datalink Linux Service Interface.

As to the libpcap library, the same explanation applies: if the NIC's exported Service Interface is consistent with Linux at the datalink layer, the library implementation would access it to provide its communication service as though it were accessing any other, functional NIC.

6. [PR] a. What Linux command displays the TCP connections currently active in a host?

netstat

b. What Linux command allows us to estimate the average Rtt to an Internet host?

ping

7. Develop the client protocol stack and the server protocol stack corresponding to an ssh connection. (The TCP port used by the server is the Well Known Port no. 22).

Consult slide titled "Internet Protocol Stack" in the presentation on Chapter 1/Section 1

8. Regarding the former question, develop now the multiplexing/encapsulation hierarchy from the standpoint of the server, *i.e.,* you must detail all the addresses and multiplexing keys involved in this communication when the server receives a new frame from the client.

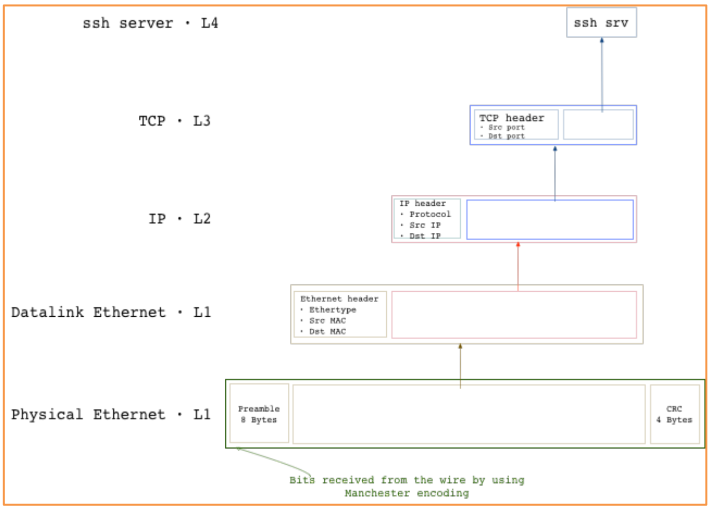

Figure 4. Protocol stack active at a receiving ssh server

9. Explain the two essential differences between the OSI and the Internet Architecture.

- OSI has 7 layers whereas TCP/IP has 4 layers
- OSI applies strict layering and TCP/IP does not

10. A data signal has a period greater than 1 μs, this signal has a bandwidth of (Tick the best option)

No solution provided to this quiz question

- ☐ a. Greater than 100 MHz
- ☐ b. Less than 100 MHz
- ☐ c. It cannot be calculated with the available data
- ☐ d. 1GHz

11. [M] An web browser connects with a web server, tick the protocols that must necessarily be activated in this connection

No solution provided to this quiz question

- ☐ a. TCP
- ☐ b. UDP
- ☐ c. http
- ☐ d. https
- ☐ e. Ethernet
- ☐ f. IP
- ☐ g. Ftp
- ☐ h. Ssh
- ☐ i. PPP
- ☐ j. X.25
- ☐ k. ssh
- ☐ l. ATM

12. Explain which TCP/IP layers are implemented by each of the following network equipment?

    a. An Internet host

    An Internet host implements all four layers

    b. A LAN switch

    The Subnetwork layer

    c. An IP router

    Subnetwork and Network(IP) layers

13. The following questions are related with the Libpcap library:

    a. Which service interface does it provide access to?

    The service interface of the Subnetwork layer

    b. The programs that use this library, do they run in kernel or in user mode? Explain why.

    User mode since they are loaded as user-space applications

    c. Why is it convenient to write network programs against this library?

    The use of this library for programming user-space network applications or protocols results in better portability across different operating systems than in the case the native RAW SOCK service interface is used.

14. [M] Tick which of the following statements about the concept of peer-to-peer interface are true

    No solution provided to this quiz question

- ☐ a. It's a term that has the same meaning as protocol
- ☐ b. It's a term that has the same meaning as service
- ☐ c. It is comprised of the legal messages exchanged by two protocol entities
- ☐ d. It is comprised of the legal messages exchanged by two service interfaces
- ☐ e. The communication medium between any two layers in two communicating hosts is any kind of physical link
- ☐ f. The communication medium between any two layers in two communicating hosts is a virtual link except in the case of the physical layer

15. Calculate the connectivity of 4 separate networks, each containing 10 nodes

    Connectivity = $4 \cdot (10 \cdot (10 - 1)) =$ **360**

16. Explain the concept of scalable connectivity. You can use a relevant example.

    The connectivity offered by a network is scalable if it preserves the communication capacity each node as more and more of them are added. This is not the case in an

Ethernet based on coaxial cable, for example, when the number of nodes exceeds some threshold. In this specific Ethernet case, what happens is that the nodes compete for the medium and resolve collisions by using CSMA/CD which, if the number of nodes is relatively high or the average level of network demand of each is high, the network bandwidth consumed in resolving collisions may turn very high.

17. Depict in a timing diagram the transmission of the bitstream 0101001110 using the NRZ, NRZ-I and Manchester line encodings

    Consult a similar example in the exam solutions published at paloalto.unileon.es/cn

18. What drawbacks about the NRZ encoding are overcome by using NRZ-i?

    Basically, NRZ-I solves the problem of long strings of 1's in the data stream to be sent. However, NRZ-I doesn't resolve the problem of long strings of bits 0 present in the data stream to be sent. In the latter case, a source encoder resolves the long strings of bits 0 by using an encoding technique known as 4B/5B. It's the output of this 4B/5B encoder that is fed into the input of the NRZ-I line encoder.

19. [M] Which of the following layer functions belong to the datalink layer

    No solution provided to this quiz question

    - ☐ a. Maintaining an acceptable level of end-to-end quality
    - ☐ b. Building frames that will encapsulate a specific type of upper-layer payload
    - ☐ c. Properly delimiting the standard fields that comprise a frame
    - ☐ d. Turning the frame's bits into signals appropriate for transmission (Line encoding)
    - ☐ e. Adding redundancy to the frame that allows the receiver to establish whether or not some error took place
    - ☐ f. Accessing the physical medium in an orderly manner that create no problems to the rest of network elements connected to it, etc.

20. The current scale of Internet is about 4000M of hosts, what protocol is responsible for this huge scale? Explain your answer.

    Basically, we claim it's the IP protocol and IP routing equipment that are responsible for this huge size. Connecting hosts by using LAN switches offers a very limited scalability due to the essential role flooding plays in this context and, also to the limited capacity of the memories used for storing MAC addresses.

21. How many networks result when a number of switches are connected?

    As more and more switches are connected to an existing network, the resulting aggregate forms a larger network, though still a single network.

22. What international institution is responsible for the TCP/IP architecture?

    IETF (Internet Engineering Task Force)